

COMP-551: Applied Machine Learning


Project #3: Modified digits

CMT submission closing Nov. 13, 11:59pm EST (=8:59pm Pacific Time).

Kaggle submission closing Nov. 13, 11:59pm EST (=4:59am UTC on next day)

Background:

For this project, you will participate in an in-class Kaggle competition on image analysis. The goal is to devise a machine learning algorithm to automatically compute mathematical functions from images. The dataset for this task is based on the classic MNIST dataset, however we have prepared a new version where each image combines two digits and a letter denoting the operation: “a” (or “A”) for “add”, and “m” (or “M”) for “multiply”. Your task is to return the result of the mathematical operation. The number of examples per class is non-uniform; you can assume that the class proportions are similar for the training and test set. Examples of the training samples are shown here:



The competition, including the data, is available here (you can use the same Kaggle account as for the previous project):

<https://www.kaggle.com/c/comp551-modified-mnist>

As for previous projects, this one should also be completed in a group of 3. Remember: you must work with different team members than for projects 1&2.

General instructions:

To participate in the competition, you must submit a list of predicted outputs for the test instances on the Kaggle website.

To solve the problem, you should try the following methods (the 3rd one is optional):

- 1) A *baseline* learner consisting of logistic regression, implemented by hand or using a library.
- 2) A *fully connected feedforward neural network* (from lecture 14), trained by backpropagation, where the network architecture (number of nodes / layers), learning rate and termination are determined by cross-validation. This method must be fully implemented by your team, and corresponding code submitted.
- 3) Any *other* machine learning method of your choice. Existing packages can be used if appropriately referenced in your report. *Suggestions*: non-linear classifier (e.g. SVM); convolutional neural networks.

For method 2, you can use algebra libraries (e.g. numpy, matlab), but the gradient descent must be implemented by you.

Your written report should provide results that compare and contrast these methods. For the Kaggle competition, you can submit results from your best performing method, from any of these categories. In addition, you can use supplementary data of your choice to enrich the training set (e.g. the original MNIST dataset); you must provide appropriate references in your report.

Submission requirements (1 submission per team, not per individual):

Same instructions as for Project 2.

Evaluation criteria:

Same criteria as for Project 2.

Submission instructions:

Same instructions as for Project 2. Submit to the Project 3 track.